The active observer gets invariant perceptions despite varying sensations. He perceives a constant object by vision despite changing sensations of light; he perceives a constant object by feel despite changing sensations of pressure; he perceives the same source of sound despite changing sensations of loudness in his ears. The hypothesis is that constant perception depends on the ability of the individual to detect the invariants, and that he ordinarily pays no attention whatever to the flux of changing sensations.

- *James Gibson* (1966, The Senses Considered as Perceptual Systems)

Chapter 5

Directional Multiscale Image Statistics

The complexities of computer vision of multivalued images suggests the development of a new class of statistically based image analysis tools. Much of this dissertation has been directed toward automatic parameter selection for nonlinear boundary preserving smoothing algorithms of multivalued images. Multivalued images are assumed to be samples of piecewise ergodic stochastic processes. In smoothing such images, the boundaries that partition the image into ergodic regions are essential elements of image structure and must be preserved. This need to identify and preserve boundary information while smoothing intensity variations resulting from noise suggests a need for directional analysis techniques. The multivalued nature of the images and the possible incommensurable nature of the multiple pixel values imply that statistics will be required.

From the piecewise ergodic assumption, it follows that the direction toward nearby image discontinuities or boundaries and the direction parallel to those boundaries are important pieces of local information. As described in section 5.1.2., differential geometry, in the form of calculus and linear algebra, provides notation and tools appropriate to the tasks of detecting the influence and orientation of nearby edges. To extract such geometric properties we seek the eigenvectors of a matrix that characterizes the structure of local image space. Given differential geometry as the model, the Hessian is the matrix of choice.

The multivalued assumption suggests that statistics, specifically central moments, be employed to achieve some commensurate measure of image properties. In statistics, the correlation matrix is a common tool for capturing and analyzing information about the interaction among related values. The result is the analysis of a local covariance matrix that captures local image properties such as direction, variances of intensity, and scale.

Chapter 4 presented a principled method for calculating local central moments of image intensity. The multiscale central moments of the previous chapter are reexamined here and supplemented with directional properties. A covariance matrix is constructed from the resulting directional moments, and an analysis of that matrix is performed. The results are eigenvalues of directional local image statistics, indicating a maximum influence of image geometry or boundaries within the image, and the minimum influence of boundaries according to the eigendirections associated with the eigenvalues.

The properties and preliminary results of directional statistics of local image intensities are presented in sections 5.1 through 5.7. Examples of local directional second order central moments are presented using two-dimensional synthetic images. These directional statistics are also generalized via canonical correlation analysis to images with two values per pixel in section 5.8. As with the isotropic multiscale central moments of the previous chapter, the resulting directional multiscale statistics are evaluated for their invariance with respect to rotation, translation, and zoom.

From a statistical viewpoint, further illumination of local geometry and the probability distribution of local intensities can be found by studying the correlations between intensity values and spatial position. This suggests treating spatial location as an inherent property of a pixel in the manner of another random variable that comprises its multivalued nature. This also can be provided by canonical correlation analysis. In Section 5.9., the correlations among the multiple intensity values of a pixel and its spatial properties are studied through canonical correlation analysis. The results show some interesting relationships between local statistics and differential geometry.

5.1. Approaches to Directional Analysis

This section covers some of the previous work on directional analysis. Other approaches to directional image analysis are briefly introduced here as a progression of ideas from orientable (or "steerable") filters to recent work based on differential geometry. The shortcomings of these approaches, namely the difficulties in generalizing their application to images of multiple incommensurable values, are addressed in section 5.1.3. These difficulties motivate the research described in the rest of the chapter.

5.1.1. Steerable Filters

Local directional image analysis has many precedents. Early examples of orientable filters exist in the literature. Simple forms of directional filters include the Kirsch and Sobel edge detection filters. Frequency based approaches to analyzing image structure through the use of directional filters include the set of Gabor filters, sinusoidal functions embedded within a Gaussian envelope. By altering the phase, frequency, and directional orientation of the sinusoid, a complete frequency decomposition of the image may be produced over a local Gaussian neighborhood. The Kirsch, Sobel, and Gabor filters are described in most introductory texts on image processing (e.g., [Jain 1985]).

Of greater interest are the scale-space representations along with their differential invariants. Chapter 2 introduced the differential structure and invariant properties of linear scale spaces generated through Gaussian and derivative-of-Gaussian filtering. A more complete discussion may be found in Florack's dissertation [Florack 1993] as well as other sources [ter Haar Romeny 1991ab]. "Steerable" directional derivatives of arbitrary order can be constructed through linear combinations of scale-space derivatives [Perona 1992]. However, this steerable property is not limited to the Gaussian as a filter function and holds for any kernel whose n-th order derivative exists [Lindeberg 1994b].

5.1.2. Matrices and Differential Geometry

The use of steerable filters often requires the computation of multiple responses of a kernel, one for every orientation or direction of interest. This requirement leads to extensive calculations requiring significant computing resources. It is tempting to reduce the computational requirement by reducing the number of orientations computed; however, this is hazardous. Undersampling the directional components of an image leads to artifacts and incorrect analysis of image structure.

These shortcomings can be overcome using differential geometry. Differential geometry supports several matrices that capture image geometry. These matrices, which include the Hessian, the second fundamental form (or shape operator), and the windowed second-moment matrix, can be analyzed for their eigenvalues and eigenvectors. This type of approach generates directional and magnitude information from a relatively small number of scale-space derivatives of the image. For example, directional curvature information of a 2D image can be extracted from the eigenanalysis of the local Hessian of an image, requiring the calculation of only four second-order derivatives at each point of interest.

The Hessian and second fundamental form have been used extensively in the analysis of the height fields of images and other 2-manifolds representing solid shape. Koenderink describes the detection of principal curvature directions and the subsequent extraction of ridges of 2-manifolds in 3-space [Koenderink 1990]. Gueziec and Ayache extract ridges of principal curvature as aids in registration of 3D datasets [Gueziec 1992]. In 2D, use of the multiscale Hessian as a differential invariant is more common. For example, Whitaker uses the Hessian in his nonlinear computations to find medial axes [Whitaker 1993ac].

Other approaches to directional analysis through differential scale-space measurement include uses of the *windowed second-moment matrix*, a linear algebraic construction that can be used in the analysis of image texture. Lindeberg uses this matrix and several measurements or statistics taken from it to perform anisotropic scale selection, junction detection, and shape from texture analysis [Lindeberg 1994a]. Weickert employs the windowed second-moment matrix in applications of nonlinear diffusion. He achieves startling results on highly figured data [Weickert 1995].

5.1.3. Intensity Invariance vs. Spatially Invariant Directional Analysis

The approaches described above share the drawbacks that they are not easily made invariant with respect to linear functions of image intensity, nor are they easily generalized to multivalued functions when the individual values cannot be considered a vector value. The multiscale central moments of the previous chapter can be applied to multiple dimensions as well as multivalued functions and retain their invariances (see section 4.9.). However, the isotropic multiscale moments do not adequately capture directional aspects and need to be generalized to include image geometry.

These difficulties suggest another approach that is invariant to transformations such as changes in contrast or gain and shifts of the background or baseline intensity. This

chapter describes the development of directional statistics using a principled system of spatial sampling. The goal is an analytic method that achieves the same qualitative results as differential analysis of scale space derivatives while preserving the invariances and multivalued properties of multiscale central moments of image intensity. Multiscale directional means and directional variances of intensity are derived from the basic definitions of joint moments. Singular value decomposition of the resulting directional covariance matrix produces eigenvalues and eigenvectors reflecting image structure. These eigenvalues have the desired invariances with respect to rotation, translation, and linear functions of intensity. Some discussion is included on generalizing these methods to multivariate images.

5.2. Directional Statistics

Consider a 2D image $I(\mathbf{p})$ with one intensity value per pixel. That is, $I(\mathbf{p})$ is defined over $\mathbb{R}^1$, where $\mathbf{p} = (p_x, p_y) \in \mathbb{R}^2$. If $I(\mathbf{p})$ is a sample of a piecewise ergodic stochastic process (Section 4.2.), its ergodicity at location $\mathbf{p}$ can be measured through multiscale central moments. Isotropic multiscale statistics are highly sensitive to fluctuations of the underlying image function (Section 4.6.). However, the geometry of the image also introduces important directional components that make directional sampling possible. For instance, an image sampled along an isointensity contour (isophote) exhibits mean ergodicity.

What is desired then is a means of sampling in the direction in which the geometry of the image contributes the least to the statistical calculation, capturing the probability distribution of the noise rather than the structure of the image. In scalar-valued images the piecewise ergodic assumption implies that the desired sampling direction is one that maximizes ergodicity; that is, the direction of least geometry is the direction that minimizes variation in local statistics. In scalar-valued images this typically means sampling in the direction of the isophote. The direction of the tangent to the isophote is a sampling direction where the image can be considered to be locally mean-ergodic. In multivalued images, the concept of isophote does not exist. The piecewise ergodic assumption can still be used to infer the direction with the least influence of image geometry. Orienting directional statistical measurements so that variation of local statistics is minimized will generate the desired direction.

Minimizing the value of directional local statistics by repetitive application of directional statistical operators across all orientations is not desirable. An alternative is to establish a matrix that captures both local geometry and local image statistics, and diagonalize it, extracting its eigenvalues and eigenvectors. Such an approach yields directional statistical analysis through a relatively small set of covariances.

The following section introduces a family of directional image covariances as the basis for a covariance matrix. The family of directional local means and directional local variances and covariances of a 2D image is presented in detail. Also the directional variances and directional covariance are shown to obey the Cauchy-Schwarz inequality that governs second-order central moments. Later sections will combine these values into matrix form for further analysis.

5.2.1. Multiscale Directional Means

How can an image be sampled along a particular locus of minimal variation? Consider the local directional mean of an image to be a Gaussian weighted sample along a line. Formally, a multiscale directional mean $\mu_{I,l}(\mathbf{p}|\sigma)$ is defined as the integral along a line *l* through the point **p** with slope $^{b}/_{a}$, where ($a^2+b^2=1$), such that

$$\mu_{I,l}(\mathbf{p}\,|\,\sigma) = \langle I(\mathbf{p});\sigma\rangle_{l} = \frac{1}{\sigma\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-\frac{\tau^2}{2\sigma^2}} I(p_x - a\tau,\, p_y - b\tau)d\tau \tag{5-1}$$

It has an associated mean $\mu_{I,\perp l}(\mathbf{p}\,|\,\sigma)$ in the conjugate direction, the direction perpendicular to the line *l*.

$$\mu_{I,\perp l}(\mathbf{p}\,|\,\sigma) = \langle I(\mathbf{p});\sigma\rangle_{\perp l} = \frac{1}{\sigma\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-\frac{\tau^2}{2\sigma^2}} I(p_x - b\tau,\, p_y + a\tau)d\tau \tag{5-2}$$

Some particular directional means of interest include those measured in the directions of the Cartesian coordinate system. The mean in the x-direction is

$$\mu_{I,x}(\mathbf{p}\,|\,\sigma) = \langle I(\mathbf{p});\sigma\rangle_{x} = G_x(\sigma,\mathbf{p}) \otimes I(\mathbf{p}) = \frac{1}{\sigma\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-\frac{\tau^2}{2\sigma^2}} I(p_x - \tau,\, p_y)d\tau \tag{5-3}$$

The y-directional mean, in the conjugate direction to the x-axis, is

$$\mu_{I,y}(\mathbf{p}\,|\,\sigma) = \langle I(\mathbf{p});\sigma\rangle_{y} = G_y(\sigma,\mathbf{p}) \otimes I(\mathbf{p}) = \frac{1}{\sigma\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-\frac{\tau^2}{2\sigma^2}} I(p_x, p_y - \tau)d\tau \tag{5-4}$$

5.2.2. Multiscale Directional Covariances

To generalize to second order directional moments, begin by considering the definitions for variance and covariance. Given two random variables $\tilde{s}$ and $\tilde{t}$, the following equations give the variance of $\tilde{s}$, $V(\tilde{s})$, the variance of $\tilde{t}$, $V(\tilde{t})$, and the covariance between $\tilde{s}$ and $\tilde{t}$, $\mathrm{Cov}(\tilde{s},\tilde{t})$, respectively.

$$\begin{aligned} V(\tilde{s}) &= \langle \tilde{s}^2 \rangle - \langle \tilde{s} \rangle^2 \\ V(\tilde{t}) &= \langle \tilde{t}^2 \rangle - \langle \tilde{t} \rangle^2 \\ \mathrm{Cov}(\tilde{s},\tilde{t}) &= \langle \tilde{s}\tilde{t} \rangle - \langle \tilde{s} \rangle \langle \tilde{t} \rangle \end{aligned} \tag{5-5}$$

Let $\tilde{s} = I_x(\mathbf{p})$, the local intensities sampled in the x-direction, and let $\tilde{t} = I_y(\mathbf{p})$, the local intensities sampled in the y-direction. What is desired is a measure of how the variation of $I(\mathbf{p})$ in the x-direction is related to the value of $I(\mathbf{p})$ in y-direction. How $I(\mathbf{p})$ varies when sampled in the x-direction or the y-direction are the two other elements of the second-order central moments.

The multiscale directional second-order central moments must reflect the same associations as shown in (5-5). That is, they should each be a difference between a

component correlating the two values (or the value with itself), and an uncorrelated term that is a product of their mean values.

Using a Gaussian as a weighting or windowing function to implement the expected value operation, the correlations between the individual values can be effected by a manipulation of the indices of integration. The variances in the x and y directions about a point **p** cannot be assumed to be uncorrelated with one another. Thus, a shared integration index in the conjugate or orthogonal direction appears in their calculations. Using these general principles, the multiscale variance in the x-axis direction becomes

$$\begin{aligned} V_{xx} &= \mu^{(2)}_{I,xx}(\mathbf{p}\,|\,\sigma) = \left\langle \left(I(\mathbf{p}) - \mu_{I,x}(\mathbf{p}\,|\,\sigma)\right)^2 \right\rangle \\ &= \frac{1}{2\pi\sigma^2}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\tau^2+\nu^2}{2\sigma^2}} \left(I(p_x-\tau, p_y-\nu) - \mu_{I,x}((p_x,p_y-\nu)\,|\,\sigma)\right)^2 d\tau\, d\nu \\ &= \frac{1}{\sigma\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-\frac{\nu^2}{2\sigma^2}} \left(\frac{1}{\sigma\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-\frac{\tau^2}{2\sigma^2}} \left(I(p_x-\tau, p_y-\nu)\right)^2 d\tau - \left(\mu_{I,x}((p_x,p_y-\nu)\,|\,\sigma)\right)^2 \right) d\nu \\ &= G(\sigma,\mathbf{p}) \otimes I(\mathbf{p})^2 - G_y(\sigma,\mathbf{p}) \otimes \left(\mu_{I,x}(\mathbf{p}\,|\,\sigma)\right)^2 \end{aligned} \tag{5-6}$$

The multiscale variance in the y-direction is similarly calculated as

$$\begin{aligned} V_{yy} &= \mu^{(2)}_{I,yy}(\mathbf{p}\,|\,\sigma) = \left\langle \left(I(\mathbf{p}) - \mu_{I,y}(\mathbf{p}\,|\,\sigma)\right)^2 \right\rangle \\ &= \frac{1}{2\pi\sigma^2}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\tau^2+\nu^2}{2\sigma^2}} \left(I(p_x-\nu, p_y-\tau) - \mu_{I,y}((p_x-\nu,p_y)\,|\,\sigma)\right)^2 d\tau\, d\nu \\ &= \frac{1}{\sigma\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-\frac{\nu^2}{2\sigma^2}} \left(\frac{1}{\sigma\sqrt{2\pi}}\int_{-\infty}^{\infty} e^{-\frac{\tau^2}{2\sigma^2}} \left((p_x-\nu, p_y-\tau)\right)^2 d\tau - \left(\mu_{I,y}((p_x-\nu,p_y)\,|\,\sigma)\right)^2 \right) d\nu \\ &= G(\sigma,\mathbf{p}) \otimes I(\mathbf{p})^2 - G_x(\sigma,\mathbf{p}) \otimes \left(\mu_{I,y}(\mathbf{p}\,|\,\sigma)\right)^2 \end{aligned} \tag{5-7}$$

For simplicity of notation, both the position parameter **p** and the scale parameter σ have been dropped from the representation V_{xx} and V_{yy}. The position and the scale or measurement aperture are always implicit in these measurements.

The local trends of the image values imply that measurements in the x and the y directions are highly correlated. Therefore, the covariance between the two must be incorporated in any descriptive image statistic. Separating the two operands and considering them separately simplifies their introduction. In the left operand of the covariance calculation, the two values are correlated as reflected in equation (5-5). In the case of directional statistics, the elements are cross-correlated in the x and y directions, sharing integration indices across their axes. Thus,

$$\left\langle (I_{x,y}(\mathbf{p}))(I_{y,x}(\mathbf{p})) \right\rangle = \frac{1}{2\pi\sigma^2}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\tau^2+\nu^2}{2\sigma^2}} \left(I(p_x-\tau, p_y-\nu)\right)\left(I(p_x-\nu, p_y-\tau)\right) d\tau d\nu \tag{5-8}$$

The right operand assumes no correlation across the two directional components. The result is a simple square of the Gaussian windowed function.

$$\begin{aligned}\left\langle (I_{x,y}(\mathbf{p}))\right\rangle\left\langle (I_{y,x}(\mathbf{p}))\right\rangle &= \left\langle I_{x,y}(\mathbf{p})\right\rangle^2 \\ &= \left(\frac{1}{2\pi\sigma^2}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\tau^2+\nu^2}{2\sigma^2}}\left(I(p_x-\tau,p_y-\nu)\right)d\tau d\nu\right)^2 \\ &= \left(\mu_I(\mathbf{p}\mid\sigma)\right)^2\end{aligned} \tag{5-9}$$

Subtracting the two operands yields the multiscale directional covariance measure.

$$\begin{aligned}V_{xy} &= \mu^{(2)}_{I,xy}(\mathbf{p}\mid\sigma) = \left\langle\left(I_{x,y}(\mathbf{p})-\mu_{I,x}(\mathbf{p}\mid\sigma)\right)\left(I_{y,x}(\mathbf{p})-\mu_{I,y}(\mathbf{p}\mid\sigma)\right)\right\rangle \\ &= \frac{1}{2\pi\sigma^2}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\tau^2+\nu^2}{2\sigma^2}}\left(I(p_x-\tau,p_y-\nu)\right)\left(I(p_x-\nu,p_y-\tau)\right)d\nu\, d\tau-\left(\mu_I(\mathbf{p}\mid\sigma)\right)^2\end{aligned} \tag{5-10}$$

5.2.3. The Cauchy-Schwarz Inequality for Multiscale Directional Covariances

The derivation of local directional central moments requires some evidence that they behave like moments of jointly distributed random variables. One condition that holds for all random variables is the Cauchy-Schwarz inequality. This same relation can be shown to exist for multiscale directional variances and covariances.

Given two random variables $\tilde{u}$ and $\tilde{v}$ the Cauchy-Schwarz inequality is stated as

$$\left\langle\tilde{u}^2\right\rangle\left\langle\tilde{v}^2\right\rangle \geq \left\langle\tilde{u}\tilde{v}\right\rangle^2 \tag{5-11}$$

Restated using directional moments, this relation for multiscale directional variances and covariances becomes:

$$V_{xx}\ V_{yy} \geq (V_{xy})^2 \tag{5-12}$$

Proof: Consider the convolution of a Gaussian kernel with the square of a sum of two continuous functions.

$$
\begin{aligned}
0 \le \frac{1}{2\pi\sigma^2}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\tau^2+\nu^2}{2\sigma^2}} & \Big(z\big(I(p_x-\tau,p_y-\nu)-\mu_{I,x}((p_x,p_y-\nu)\,|\,\sigma)\big) \\
& + \big(I(p_x-\nu,p_y-\tau)-\mu_{I,y}((p_x-\nu,p_y)\,|\,\sigma)\big)\Big)^2 d\tau\, d\nu \\
= z^2\Bigg(\frac{1}{2\pi\sigma^2}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\tau^2+\nu^2}{2\sigma^2}} & \Big(\big(I(p_x-\tau,p_y-\nu)-\mu_{I,x}((p_x,p_y-\nu)\,|\,\sigma)\big)^2\Big) d\tau\, d\nu\Bigg) \\
+ 2z\Bigg(\frac{1}{2\pi\sigma^2}\Bigg(\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\tau^2+\nu^2}{2\sigma^2}} & \Big(\big(I(p_x-\tau,p_y-\nu)-\mu_{I,x}((p_x,p_y-\nu)\,|\,\sigma)\big) \\
& \big(I(p_x-\nu,p_y-\tau)-\mu_{I,y}((p_x-\nu,p_y)\,|\,\sigma)\big)\Big) d\tau\, d\nu\Bigg)\Bigg) \\
+ \frac{1}{2\pi\sigma^2}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{\tau^2+\nu^2}{2\sigma^2}} & \Big(\big(I(p_x-\nu,p_y-\tau)-\mu_{I,y}((p_x-\nu,p_y)\,|\,\sigma)\big)^2\Big) d\tau\, d\nu \\
= z^2 V_{xx} + 2zV_{xy} + V_{yy} &
\end{aligned}
\tag{5-13}
$$

Equation (5-13) can be rewritten as

$$z^2 V_{xx} + 2zV_{xy} + V_{yy} \ge 0 \tag{5-14}$$

Inserting a non-negative constant $d \ge 0$ modifies (5-14) to a homogeneous quadratic form.

$$z^2 V_{xx} + 2zV_{xy} + V_{yy} - d = 0 \quad , \quad \text{where } d \ge 0 \tag{5-15}$$

Invoking the quadratic formula yields

$$z = \frac{-V_{xy} \pm \sqrt{4V_{xy}^{\,2} - 4(V_{xx}(V_{yy}-d))}}{2V_{xx}} \quad , \quad \text{where } d \ge 0 \tag{5-16}$$

Equations (5-14) and (5-16) imply the following:

- if $4V_{xy}^{\,2} - 4(V_{xx}(V_{yy}-d)) < 0$, then both roots are imaginary.
- if $4V_{xy}^{\,2} - 4(V_{xx}(V_{yy}-d)) = 0$, then there is only one real double root.
- $4V_{xy}^{\,2} - 4(V_{xx}(V_{yy}-d))$ cannot be greater that zero, implying two real roots; this condition would require the curve of the original function in (5-14) to pass across zero in two places, making it negative across some interval of z.

The requirement for $4V_{xy}^{\,2} - 4(V_{xx}(V_{yy}-d)) \le 0$ implies the following relationships.

$$
\begin{aligned}
0 &\geq 4V_{xy}^{\,2} - 4(V_{xx}(V_{yy} - d)) \\
0 &\geq V_{xy}^{\,2} - V_{xx}V_{yy} + dV_{xx} \\
V_{xx}V_{yy} &\geq V_{xy}^{\,2} + dV_{xx}
\end{aligned}
\tag{5-17}
$$

Since d 0, by transitivity the inequality implies the following relations:

$$
V_{xx}V_{yy} \geq dV_{xx} \quad \text{and} \quad V_{xx}V_{yy} \geq V_{xy}^{\,2} \tag{5-18}
$$

❑.

5.3. Directional Multiscale Statistics of Sample 2D Images

The statistics V_{xx}, V_{xy}, and V_{yy} described above are easily computed on 2D image data. Consider the image of Figure 5.1. This computer generated image I(**p**) contains a tapered test object with straight and curved boundaries. White noise $\tilde{u}$ (Gaussian distributed, zero-mean, spatially-uncorrelated) has been added. The resulting image, $\tilde{I}(\mathbf{p}) = I(\mathbf{p}) + \tilde{u}$, is shown below.

As in earlier chapters, the term signal to noise ratio (SNR) when applied to the examples in this chapter will refer to the difference of the foreground intensity and the background intensity divided by the standard deviation of the additive spatially uncorrelated noise. The pixel is the atomic image element. Thus, the additive noise and consequently the relative measurement of noise to signal is expressed as the SNR per pixel. The measured SNR per pixel in Figure 5.1 has been set to 4:1 on a raster resolution of 256 x 256 pixels.

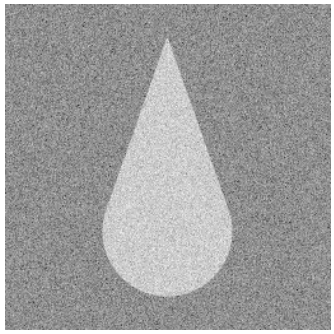

Figure 5.1. A test image with SNR of 4:1 with a raster resolution of 256×256 pixels.

The images of Figure 5.2 are three local directional second order moments of Figure 5.1 using an aperture whose scale or spatial aperture is 2 pixels wide. Figure 5.2a shows $\mu^{(2)}_{\tilde{I},xx}(\mathbf{p}\,|\,\sigma)$, the variance as measured in the x-direction. Figure 5.2b shows $\mu^{(2)}_{\tilde{I},xy}(\mathbf{p}|\sigma)$, the covariance between the x-directional and y-directional samples. Figure 5.2c shows $\mu^{(2)}_{\tilde{I},yy}(\mathbf{p}\,|\,\sigma)$, the variance as measured in the y-direction. A combination of all three moments and joint moments are required to capture the information of the second-order directional moments.

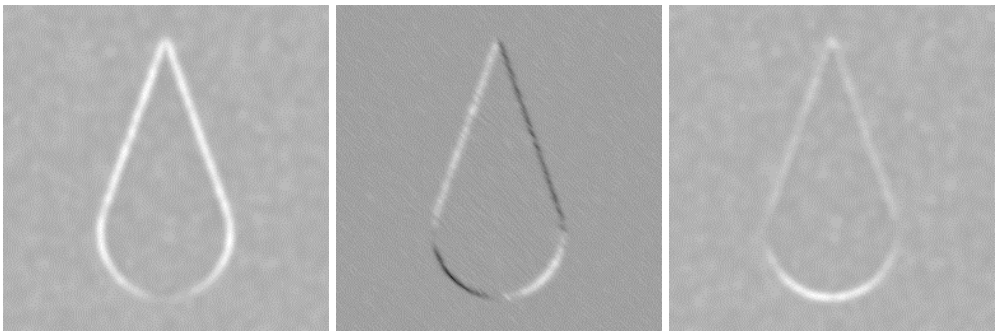

a. b. c.

Figure 5.2. Directional variances of the objects from Figure 5.1. (a: $V_{xx} = \mu^{(2)}_{\tilde{I},xx}(\mathbf{p}\,|\,\sigma)$, b: $V_{xy} = \mu^{(2)}_{\tilde{I},xy}(\mathbf{p}|\sigma)$, c: $V_{yy} = \mu^{(2)}_{\tilde{I},yy}(\mathbf{p}\,|\,\sigma)$). In all images, σ = 2 pixels and grey represents a value of 0. Bright grey to white indicates positive values, and dark grey to black indicates negative values. The origin is in the upper left corner, the x-axis oriented to the right, and the y-axis oriented from the toward the bottom of the page.

These directional covariances reflect local image geometry as well as characterize the noise within the image. Figure 5.2a, the variance in the x-direction $\mu^{(2)}_{\tilde{I},xx}(\mathbf{p}\,|\,\sigma)$, shows a relatively constant measure in the background and in the center of the teardrop figure. These values, when queried, are approximately the value for the variance of the additive noise in the image. There is a strong effect around the perimeter of the teardrop figure. Where $\mu^{(2)}_{\tilde{I},xx}(\mathbf{p}\,|\,\sigma)$ is sampling across the boundary between the object and background, a bimodal distribution of image intensities is observed, resulting in a large local directional variance measure. This response peaks at the two sides of the teardrop where the tangent to the object boundary is perpendicular to the sampling direction. At the base of the teardrop, the x-direction (the sampling direction) and the object boundary are parallel. At this point, a single mode is detected in the directional sample of image intensities, and the variance shows the same value as the measure of background noise.

Correspondingly, in Figure 5.2c, when measuring $\mu^{(2)}_{\tilde{I},yy}(\mathbf{p}\,|\,\sigma)$ the local variance in the y-direction, there are two such points on opposite sides of the teardrop where the tangent to the perimeter of the figure is parallel to the y-axis. The maximum response of the variance sampled in the y-direction occurs at the base of the teardrop where the sampling direction is perpendicular to the object boundary creating a bimodal distribution of image intensities in the directional sample. In locally ergodic regions, there is a constant response approximately equal to the variance of the additive noise. The greater directional variances at the sloping sides of the teardrop in Figure 5.2a as compared to that in Figure 5.2c corresponds to the fact that the orientation of the boundary is more vertical than horizontal. The response along the upper boundary in Figure 5.2c is lower than the variance shown in Figure 5.2a, implying that the boundary is closer to being parallel with the y-direction than the x-direction.

Figure 5.2b shows the covariance measure $\mu^{(2)}_{\tilde{I},xy}(\mathbf{p}|\sigma)$. As expected, in places where a discontinuous change in intensities are correlated with the x and y Cartesian directions, a

positive response is detected. Those areas where boundaries are negatively correlated between the x and y directions show a negative response. In locally ergodic regions (the center of the teardrop and the perimeter of the image), $\mu^{(2)}_{\tilde{I},xy}(\mathbf{p}|\sigma)$, like the other directional variance measures, also computes a constant value approximately equal to the variance of the additive spatially uncorrelated noise in the image.

The three scaled statistical representations of Figure 5.2 are significant for image processing tasks. The figures 5.2a and 5.2c have responses similar to the square of the partial derivatives of I(**p**) in the x-direction and the y-direction. The covariance image 5.2b reflects the relationship between the two Cartesian directions and has a response similar to the product of the partial derivatives in the x and y directions.

5.4. The Directional Multiscale Covariance Matrix

Having defined and justified local second-order directional central moments of image intensity, the next step is to organize them into matrix form for further analysis. This section describes the local covariance matrix. It also presents the analysis of the general form of the covariance matrix through singular value decomposition (SVD). The resulting equations for the eigenvalues and eigenvectors of the covariance matrix are described here.

Given the three values V_{xx}, V_{yy}, and V_{xy}, it is straightforward to construct the covariance matrix $\mathbf{C}_{I(\mathbf{p})}$, which describes how the behavior of I over x and the behavior of I over y vary or covary.

$$\mathbf{C}_{I(\mathbf{p})} = \begin{bmatrix} V_{xx} & V_{xy} \\ V_{xy} & V_{yy} \end{bmatrix} \tag{5-19}$$

As with the variance notation of V_{xx}, V_{yy}, and V_{xy}, the scale parameter has been dropped from the representation $\mathbf{C}_{I(\mathbf{p})}$ to simplify the notation. Scale as a parameter in these measurements is assumed.

The Cauchy-Schwarz inequality implies that the determinant of the covariance matrix is non-negative. A positive determinant ensures that the covariance matrix is of full rank.

If its rank is full, the covariance matrix $\mathbf{C}_{I(\mathbf{p})}$ for point **p** can be analyzed through an eigenanalysis, determining its eigenvalues and eigenvectors. This process of singular value decomposition yields some important statistics that have significant invariant properties. Pertinent linear algebraic properties are provided in the appendix of this chapter. Their application to the covariance matrix is summarized here.

The eigenvalues of $\mathbf{C}_{I(\mathbf{p})}$ are

$$\lambda_1 = \frac{(V_{xx} + V_{yy}) + \sqrt{(V_{xx} + V_{yy})^2 - 4(V_{xx}V_{yy} - V_{xy}{}^2)}}{2} \tag{5-20a}$$

$$\lambda_2 = \frac{(V_{xx} + V_{yy}) - \sqrt{(V_{xx} + V_{yy})^2 - 4(V_{xx}V_{yy} - V_{xy}^2)}}{2} \tag{5-20b}$$

The diagonal symmetry of the covariance matrix $\mathbf{C}_{I(\mathbf{p})}$ indicates that its eigenvalues are both real and non-negative. If $V_{xx}V_{yy} > V_{xy}^2$, then the eigenvalues are both positive and real, so the covariance matrix is positive definite. If $V_{xx}V_{yy} = V_{xy}^2$, the rank of the covariance matrix is not full, and there is a single eigenvalue representing isotropically distributed variance.

The eigenvalues of the local covariance matrix λ_1 and λ_2 are principal values revealing information regarding the shape and structure of local image geometry in much the same way as the eigenvalues of the Hessian describe the intensity surface. λ_1 and λ_2 are invariant with respect to image rotation, translation and zoom (simultaneous multiplication of image resolution and measurement aperture).

The eigenvectors **u** and **v** corresponding to λ_1 and λ_2, respectively, are

$$\mathbf{u} = \left(\frac{(\lambda_1 - V_{yy})}{\sqrt{V_{xy}^2 + (\lambda_1 - V_{yy})^2}} \quad , \quad \frac{V_{xy}}{\sqrt{V_{xy}^2 + (\lambda_1 - V_{yy})^2}} \right)$$

and (5-21)

$$\mathbf{v} = \left(\frac{(\lambda_2 - V_{yy})}{\sqrt{V_{xy}^2 + (\lambda_2 - V_{yy})^2}} \quad , \quad \frac{V_{xy}}{\sqrt{V_{xy}^2 + (\lambda_2 - V_{yy})^2}} \right)$$

These vectors are unit vectors signifying the principal variance directions. The eigenvalues reflect the magnitudes of the expression of the directional variance in local image space. The vector **u** is the direction of maximum local variance at **p**, while **v** is the minimum variance direction. These vectors are orthogonal.

The eigenvectors comprise the diagonalizing matrix **K** = (**u**, **v**). This matrix is a linear transformation, a rotation in this case, that aligns the sampling directions along the **u** and **v** directions. **K** diagonalizes the covariance matrix $\mathbf{C}_{I(\mathbf{p})}$ in the following fashion.

$$\mathbf{K}\, \mathbf{C}_{I(\mathbf{p})}\, \mathbf{K}^T = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \tag{5-22}$$

The **v** direction is called the *ergodic direction*, indicating the direction in which the probability distribution of intensities shows the greatest ergodicity. λ_2 is the variance sampled in the ergodic direction **v**, and represents the local directional variance measurement with the least influence of image geometry. This implies that λ_2 is a reasonable measure of the variance of the local noise process since the contribution of local image structure has been minimized. λ_2 can therefore be used to normalize directional statistics, enabling measurements that are invariant with respect to linear functions of intensity.

In those locations where the distribution is isotropic with respect to space (i.e., there is no preferred direction), λ_1 approximately equals λ_2, the **v** direction becomes ambiguous, and isotropic sampling methods discussed in Chapter 4 are applicable.

5.5. SVD Applied to Directional Multiscale Statistics of 2D Images

When applied to local covariance measures of 2D images, singular value decomposition should generate, as eigenvalues of the directional covariance matrix, two variance values, one showing the maximum influence of boundaries and other geometry within the image and one local variance value that minimizes the influence of nearby boundaries. Essentially, these variance values are the result of linear transformations of the covariance matrix that indicated the directions of least and greatest impact of image geometry on directional variances.

The analysis described in section 5.4. is easily applied to 2D image data. Consider the source image of Figure 5.1. The image values for the directional variances and covariance shown in Figure 5.2 can be simplified through singular value decomposition to generate eigenvalue images. Figure 5.3 shows the major and minor eigenvalues for the image in Figure 5.1, measured with an aperture of 2 pixels.

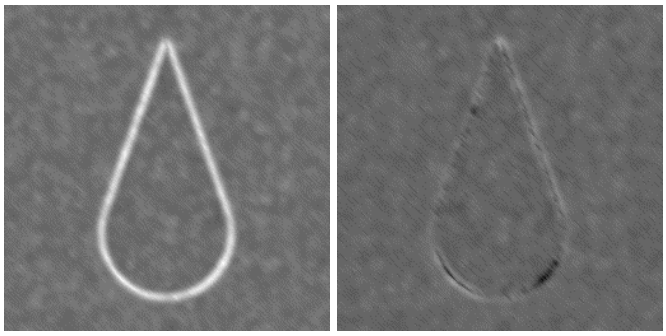

a. b.

Figure 5.3. Eigenvalue images of the object from Figure 5.1, computed with a spatial aperture σ of 2 pixels. (a: λ_1, b: λ_2). In both of these images, black is zero and bright indicates positive values.

The eigenvalues are directional variances that have been subjected to linear transformations that maximize and minimize the influence of image geometry. The λ_1 values in Figure 5.3a have a relatively large constant response around the object boundary. The fluctuations of directional variances in the Cartesian x and y directions from Figure 5.2 are no longer apparent. The λ_2 in figure 5.3b shows a relatively constant value across the image. When queried, the values of the λ_2 image are centered about the variance of the additive noise of the image. There are numerical artifacts and effects of isophote curvature that prevent a complete constant response in the λ_2 image, but the predominant effect is a constant evaluation of image variance that minimizes the influence of first order elements of image structure.

Figure 5.4 shows the vector field associated with the eigenvector **u** corresponding to the larger eigenvalue λ_1. These vectors reflect the direction in which there is maximum variance at each pixel. The lengths of the arrows show relative magnitudes of the λ_1 eigenvalues. Near the boundary of the teardrop, the directions of these vectors, as expected, are approximately orthogonal to the boundary. The set of **v** eigenvectors, corresponding to the smaller of the eigenvalues λ_2, are perpendicular to the eigenvectors

shown. Near the boundary of the teardrop the **v** eigenvectors are approximately parallel to the boundary.

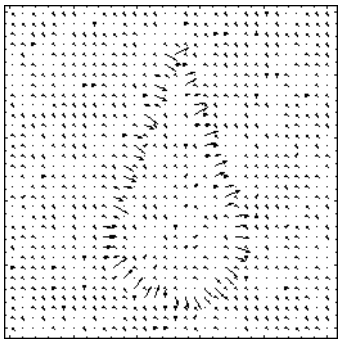

Figure 5.4. Eigenvector image of the object from Figure 5.1, computed with a spatial aperture σ of 2 pixels. The image reflects only the eigenvector **u** in the direction of maximum variance at each pixel; the eigenvector **v** in the direction of minimum variance is perpendicular to the vectors shown. The lengths of the vector representations indicate relative magnitude, reflecting the maximum variance or eigenvalue.

5.6. Multiscale Gauge Coordinates of Image Statistics

In their description of differential scale space representations, ter Haar Romeny, Florack, and Lindeberg each make use of *gauge coordinates*, a local coordinate frame formed by the isophote tangent and gradient direction describing a natural orthogonal basis at each location within the image (see Chapter 2, [ter Haar Romeny 1991ab], [Florack 1993], and [Lindeberg 1994b]). When scale space differential invariants are recast in gauge coordinates, the natural coordinate system creates a framework for easier interpretation. This simplification is reflected in a reduction in the complexity of the notation.

Eigenanalysis of directional multiscale covariances generates a similar coordinate frame hereafter called the *covariance gauge*. As previously mentioned, the eigenvectors represent principal variance directions, are orthogonal, and make a natural coordinate frame for analyzing local variations of intensity. Notation in the covariance gauge coordinate system is greatly simplified over the pixel grid directions. The eigenvalues are defined to be the variance in the **u** and **v** directions, respectively.

$$V_{uu} = \frac{(V_{xx} + V_{yy}) + \sqrt{(V_{xx} + V_{yy})^2 - 4(V_{xx} V_{yy} - V_{xy}{}^2)}}{2} \tag{5-23a}$$

$$V_{vv} = \frac{(V_{xx} + V_{yy}) - \sqrt{(V_{xx} + V_{yy})^2 - 4(V_{xx} V_{yy} - V_{xy}{}^2)}}{2} \tag{5-23b}$$

$$V_{uv} = 0 \tag{5-23c}$$

The magnitudes of the gauge, measured through the eigenvalues of the singular value decomposition provide a metric by which other geometric invariants may be normalized. For example, scale-space gradients may be normalized by the square root of the minimum

variance at a pixel location, making it invariant with respect to linear functions of intensity. The eigenvalues of the covariance matrix and the gradient magnitude increase and decrease proportionally under transformations of linear functions of intensity. Their quotient is an invariant with respect to linear shifts in intensity.

The covariance gauge exists only if the *covariance gauge conditions* are met. The derivative based gauge coordinates of ter Haar Romeny and Florack exist only if the image intensity gradient exists and is unique [ter Haar Romeny 1991ab]. The covariance gauge exists if there are distinct principal variance directions. That is, it exists if unique eigenvectors can be found for the covariance matrix. This condition, the property that the covariance matrix is diagonalizable, is called the covariance gauge condition and is a generic property of images.

5.7. Invariants of Directional Multiscale Moments

A stated goal of the development of local directional central moments was that their measurement be made invariant with respect to spatial rotation, translations, and zoom (the combined magnification or minification of measurement aperture and image or pixel scale), as well as invariant with respect to linear functions of intensity. Up to this point this chapter has been a progression of steps demonstrating local directional statistics of image intensity with each step incrementally showing additional invariances of the second-order local directional central moments of intensity. The local directional covariance matrix is invariant with respect to spatial translation, and after rotation to the covariance gauge, it is invariant to rotation. However, it is not invariant with respect to linear functions of intensity. This section describes invariant measures of local image statistics that have all the desired invariances with respect to changes in space as well as to linear changes in intensity.

Local image measurements that are invariant with respect to linear functions of intensity have been calculated before using quantities from differential geometry, such as the windowed second moment matrix. Lindeberg uses the windowed second moment matrix to measure and approximate adjustments to his scale parameters, allowing him to steer his sampling aperture. The result is that the size and orientation of his anisotropic Gaussian sampling aperture becomes a function of image intensity, based on anisotropy measurements made from the windowed second moment matrix [Lindeberg 1994a]. Using these techniques, he achieves some startling results in shape from texture algorithms.

While the covariance matrix $\mathbf{C}_{\mathrm{I}(\mathbf{p})}$ derived in Chapter 5 is distinct from the windowed second moment matrix prominent in Lindeberg's shape from texture analysis and Weickert's texture based nonlinear diffusion [Weickert 1995], it shares many of the same properties. In particular, some of the same invariants presented by Lindeberg regarding the windowed second moment matrix apply to the directional multiscale covariance matrix. Unlike Lindeberg's local second moment matrix, no adjustments for anisotropic Gaussian sampling kernels are necessary. The difficulty that arises due to the lack of independence between the zeroth order sampling aperture and the first derivative that exists in the windowed second moment matrix does not apply to the local directional

covariance matrix. A single scale parameter can be used in the local covariance matrix $\mathbf{C}_{I(\mathbf{p})}$ simplifying the construction of statistical invariants.

Adapting from Lindeberg [Lindeberg 1994a], consider the following image descriptors.

$$P = V_{xx} + V_{yy}, \quad C = V_{xx} - V_{yy}, \quad \text{and} \quad S = 2\,V_{xy} \tag{5-24}$$

$P = \text{Trace}(\mathbf{C}_{I(\mathbf{p})})$ can be interpreted as a measure of the strength of the variance response. The other two descriptors C and S contain directional information, which can be summarized into $\hat{Q}$, a normalized measure of local anisotropy.

$$Q = \sqrt{C^2 + S^2} \quad \text{and} \quad \hat{Q} = \frac{Q}{P} \tag{5-25}$$

$Q = \sqrt{(V_{xx} + V_{yy})^2 - 4(V_{xx}V_{yy} - V_{xy}^2)}$ is the discriminant of the eigenvalue equation. As indicated in section 5.4, if Q = 0, there is only one eigenvalue and no preferred eigendirection, implying a isotropic distribution of image intensities. Q is therefore interpreted as a measure of local image anisotropy. Q is normalized by $\text{Trace}(\mathbf{C}_{I(\mathbf{p})})$ to generate a dimensionless statistic. Clearly, for the directional covariance based measure of *normalized anisotropy*, $\hat{Q} \in [0,1]$ (if the Cauchy-Schwarz assumption holds), and

1. $\hat{Q} = 0$ if and only if $V_{xx} = V_{yy}$ and $V_{xy} = 0$.
2. $\hat{Q} = 1$ if and only if $V_{xx}V_{yy} = V_{xy}^2$.

Note that the $\hat{Q}$ statistic is invariant with respect to spatial translation and with respect to linear functions of intensity.

Figure 5.5a is a test image where the figures in the foreground and the background pixels have the same mean intensity (0) and the same variance (6.25). The difference between the two region types is that the foreground has a directional component not necessarily one of the Cartesian directions. Figure 5.5b is the test image processed for anisotropy. The normalized value of $\hat{Q}$ is portrayed with the values ranging from 0.003 in the background to 0.89 in the foreground.

Values for the $\hat{Q}$ that are close to 1 indicate strong local anisotropy. Values near 0 indicate an isotropic distribution of noise and geometry. Both of these effects are seen in the example above. The measurements reflected in Figure 5.5b show that within the object boundaries, strong anisotropy is detected with $\hat{Q}$ near 0.9. The background measures in Figure 5.5b are near 0 reflecting no directional preference, an isotropic distribution of intensities.

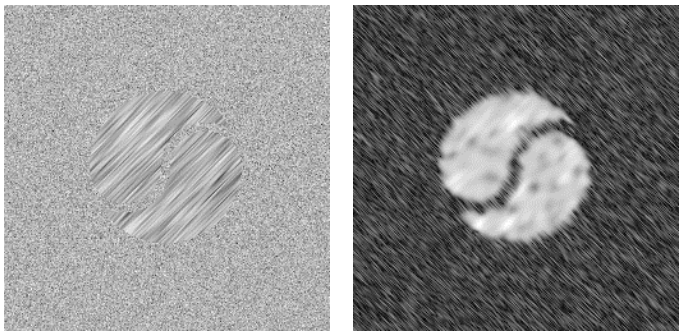

a. b.

Figure 5.5. (a) Test figure exhibiting significant directional spatial correlation and (b) the local anisotropy statistic $\hat{Q}$ where $\sigma = 3$. In both images, the raster resolution is 256×256.

5.8. Multiscale Directional Statistics of Multivalued Images

A continuing feature of this research has been the generalizing of scalar techniques for use on multivalued image applications. Directional statistics are expected to have their greatest impact when applied to multivalued images because principled tools for directional analysis of multiple incommensurable values per pixel via differential geometric operators are not available. The premise of this work is that images with incommensurable values within each pixel can be studied by exploring the local intensity distributions through multiscale central moments (Section 4.6) and then making the various intensity values at a pixel commensurable via their observed covariances. To compensate for spatial biases in ergodicity introduced by local image geometry, the eigenvalues of the multivalued directional local covariance matrix, or the equivalent calculated correlations, are required. As described in the following section, these values correlate the magnitudes of the directional covariances of the local intensity values. The geometry of these correlations, their trends and changes over space, can be used to analyze image structure even when the original values of the image are incommensurable.

Recall that in the scalar-valued case, the λ_2 statistic characterizes the probability distribution of local noise. The directional covariance matrix of an image function I(**p**) is calculated and simplified using singular value decomposition (or SVD). The covariance matrix $\mathbf{C}_{\mathrm{I}(\mathbf{p})}$ is rotated into a coordinate frame designated by direction vectors **u** and **v** in the linear operator **K**.

$$\mathbf{K}\,\mathbf{C}_{\mathrm{I}(\mathbf{p})}\,\mathbf{K}^{\mathrm{T}} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \tag{5-26}$$

K is oriented to maximize $\lambda_1 = \mu^{(2)}_{\mathrm{I,uu}}(\mathbf{p}|\sigma)$, the variance sampled in the **u** direction. The **v** direction is orthogonal to **u** and indicates the sampling direction of minimum variance, which I call the *ergodic direction*. Near object boundaries, **v** is the direction in which the probability distribution of intensities shows the greatest ergodicity. $\lambda_2 = \mu^{(2)}_{\mathrm{I,vv}}(\mathbf{p}|\sigma)$ is the variance sampled in the ergodic direction **v**. λ_2 represents the local

directional variance measurement with the least influence of image geometry. This implies that λ_2 is a reasonable measure of the variance of the local noise process.

The methods of singular value decomposition are insufficient when the components of the covariance matrix are not scalar quantities. The calculation of correlations among multivalued directional statistics requires the use of *canonical correlation analysis*. While SVD will reduce a matrix of scalar values, canonical correlation analysis can be applied to partitioned or block matrices (a matrix of matrices or other tensors). Canonical correlation analysis can therefore be used to analyze random variables containing multiple values each, while singular value decomposition is used to calculate covariances among random variables with scalar components. The two methods are related and can be shown to generate identical results when analyzing covariances of scalar-valued random variables.

5.8.1. Canonical Correlation Analysis of Multivalued Directional Statistics

Canonical correlation analysis is a statistical approach for simplifying a symmetric matrix to its principal components. This section applies canonical correlation analysis to multivalued directional statistics. This use of canonical correlation analysis is a simplified adaptation from Arnold's description of these techniques [Arnold 1981].

Consider the case of a 2D two-valued image. Let $\mathbf{I}(\mathbf{p})$ be defined over $\mathbb{R}^2$ such that $\mathbf{I}(\mathbf{p}) = (I_1(\mathbf{p}), I_2(\mathbf{p}))$. Define $\mathbf{p} = (p_x, p_y) \in \mathbb{R}^2$. Let the $I_1(\mathbf{p})$ values sampled in the x-direction and the $I_1(\mathbf{p})$ values sampled in the y-direction be considered as separate elements of a two-valued random variable, $(I_{1,x}(\mathbf{p}), I_{1,y}(\mathbf{p}))$. Similarly, let the $I_2(\mathbf{p})$ values sampled in the x-direction and the $I_2(\mathbf{p})$ values sampled in the y-direction be considered as separate elements of a two-valued random variable, $(I_{2,x}(\mathbf{p}), I_{2,y}(\mathbf{p}))$. The analysis of each of these multivalued random variable pairs requires the calculation of a linear transformation or rotation $\mathbf{G_I}$ (where $\mathbf{I} = (I_1, I_2)$) that decorrelates $\mathbf{I}_x(\mathbf{p})$ and $\mathbf{I}_y(\mathbf{p})$. Two matrices $\mathbf{G}_1$ and $\mathbf{G}_2$ comprise $\mathbf{G_I}$. $\mathbf{G}_1$ and $\mathbf{G}_2$ are closely related to their counterpart $\mathbf{K}$ in the scalar calculations of singular value decomposition shown above. However, $\mathbf{G}_1$ and $\mathbf{G}_2$ not only decorrelate the directional variations within-variable elements, but by applying them to the cross covariance matrix, they also decorrelate the two-valued directional variances across their multiple elements.

Let Σ be the covariance matrix of $(I_{1,x}(\mathbf{p}), I_{1,y}(\mathbf{p}))$ and $(I_{2,x}(\mathbf{p}), I_{2,y}(\mathbf{p}))$. That is,

$$\Sigma = \begin{pmatrix} \Sigma_{I_1 I_1} & \Sigma_{I_1 I_2} \\ \Sigma_{I_2 I_1} & \Sigma_{I_2 I_2} \end{pmatrix} \tag{5-27}$$

Σ_{I_1,I_1}, Σ_{I_1,I_2}, Σ_{I_2,I_1}, and Σ_{I_2,I_2} are all 2×2 matrices, with $\Sigma_{I_2,I_1} = \Sigma_{I_1,I_2}{}^{T}$. Σ_{I_1,I_1} and Σ_{I_2,I_2} are the familiar directional covariance matrices applied to the separate intensity values.

$$\Sigma_{I_1,I_1} = \mathbf{C}_{I_1(\mathbf{p})} = \begin{pmatrix} V_{I_1,xx} & V_{I_1,xy} \\ V_{I_1,xy} & V_{I_1,yy} \end{pmatrix} = \begin{pmatrix} \mu^{(2)}_{I_1,xx}(\mathbf{p}|\sigma) & \mu^{(2)}_{I_1,xy}(\mathbf{p}|\sigma) \\ \mu^{(2)}_{I_1,xy}(\mathbf{p}|\sigma) & \mu^{(2)}_{I_1,yy}(\mathbf{p}|\sigma) \end{pmatrix} \tag{5-28a}$$

$$\Sigma_{I_2,I_2} = \mathbf{C}_{I_2(\mathbf{p})} = \begin{pmatrix} V_{I_2,xx} & V_{I_2,xy} \\ V_{I_2,xy} & V_{I_2,yy} \end{pmatrix} = \begin{pmatrix} \mu^{(2)}_{I_2,xx}(\mathbf{p}|\sigma) & \mu^{(2)}_{I_2,xy}(\mathbf{p}|\sigma) \\ \mu^{(2)}_{I_2,xy}(\mathbf{p}|\sigma) & \mu^{(2)}_{I_2,yy}(\mathbf{p}|\sigma) \end{pmatrix} \tag{5-28b}$$

The covariance matrix between the two image intensity values and their directional components Σ_{I_1,I_2} is 2×2 and corresponds to

$$\Sigma_{I_1,I_2} = \begin{pmatrix} \mu^{(2)}_{I_1,x;I_2,x}(\mathbf{p}|\sigma) & \mu^{(2)}_{I_1,x;I_2,y}(\mathbf{p}|\sigma) \\ \mu^{(2)}_{I_1,y;I_2,x}(\mathbf{p}|\sigma) & \mu^{(2)}_{I_1,y;I_2,y}(\mathbf{p}|\sigma) \end{pmatrix} \tag{5-29}$$

Assuming that Σ_{I_1,I_1} and Σ_{I_2,I_2} are reducible via singular value decomposition, there exist $\mathbf{K}_1$ and $\mathbf{K}_2$, invertible 2×2 diagonalizing matrices such that

$$\mathbf{K}_1 \Sigma_{I_1,I_1} \mathbf{K}_1^{\mathrm{T}} = \begin{pmatrix} \lambda_{1,I_1} & 0 \\ 0 & \lambda_{2,I_1} \end{pmatrix} = \begin{pmatrix} \mu^{(2)}_{I_1,\mathbf{u}_1\mathbf{u}_1}(\mathbf{p}|\sigma) & 0 \\ 0 & \mu^{(2)}_{I_1,\mathbf{v}_1\mathbf{v}_1}(\mathbf{p}|\sigma) \end{pmatrix}$$

and (5-30)

$$\mathbf{K}_2 \Sigma_{I_2,I_2} \mathbf{K}_2^{\mathrm{T}} = \begin{pmatrix} \lambda_{1,I_2} & 0 \\ 0 & \lambda_{2,I_2} \end{pmatrix} = \begin{pmatrix} \mu^{(2)}_{I_2,\mathbf{u}_2\mathbf{u}_2}(\mathbf{p}|\sigma) & 0 \\ 0 & \mu^{(2)}_{I_2,\mathbf{v}_2\mathbf{v}_2}(\mathbf{p}|\sigma) \end{pmatrix}$$

As in the case of scalar-valued images, $\mathbf{K}_1$ rotates Σ_{I_1,I_1} onto the $\mathbf{u}_1$-$\mathbf{v}_1$ coordinate frame. $\mathbf{K}_1$ is oriented to maximize the eigenvalue $\lambda_{1,I_1} = \mu^{(2)}_{I_1,\mathbf{u}_1\mathbf{u}_1}(\mathbf{p}|\sigma)$, the variance sampled in the $\mathbf{u}_1$ direction at $I_1(\mathbf{p})$. $\mathbf{v}_1$ is orthogonal to $\mathbf{u}_1$, and the eigenvalue $\lambda_{2,I_1} = \mu^{(2)}_{I_1,\mathbf{v}_1\mathbf{v}_1}(\mathbf{p}|\sigma)$, the variance sampled along the $\mathbf{v}_1$ direction is the minimum variance with respect to orientation. Likewise, $\mathbf{K}_2$ rotates Σ_{I_2,I_2} onto the $\mathbf{u}_2$-$\mathbf{v}_2$ coordinate frame with eigenvalues λ_{1,I_2} and λ_{2,I_2} reflecting maximum and minimum variance measurements with respect to changes in orientation.

Choose $\mathbf{G}_1$ and $\mathbf{G}_2$ such that

$$\mathbf{G}_1 = \begin{pmatrix} \lambda_{1,I_1}^{-1/2} & 0 \\ 0 & \lambda_{2,I_1}^{-1/2} \end{pmatrix} \mathbf{K}_1 \qquad \mathbf{G}_1^{\mathrm{T}} = \mathbf{K}_1^{\mathrm{T}} \begin{pmatrix} \lambda_{1,I_1}^{-1/2} & 0 \\ 0 & \lambda_{2,I_1}^{-1/2} \end{pmatrix}$$

and (5-31)

$$\mathbf{G}_2 = \begin{pmatrix} \lambda_{1,I_2}^{-1/2} & 0 \\ 0 & \lambda_{2,I_2}^{-1/2} \end{pmatrix} \mathbf{K}_2 \qquad \mathbf{G}_2^{\mathrm{T}} = \mathbf{K}_2^{\mathrm{T}} \begin{pmatrix} \lambda_{1,I_2}^{-1/2} & 0 \\ 0 & \lambda_{2,I_2}^{-1/2} \end{pmatrix}$$

$\mathbf{G}_1$ and $\mathbf{G}_2$ diagonalize and normalize the individual directional covariance matrices.

$$\mathbf{G}_1 \Sigma_{I_1,I_1} \mathbf{G}_1^{\mathrm{T}} = \mathbf{I} \qquad \mathbf{G}_2 \Sigma_{I_2,I_2} \mathbf{G}_2^{\mathrm{T}} = \mathbf{I} \tag{5-32}$$

$\mathbf{G}_1$ and $\mathbf{G}_2$ diagonalize the directional covariance matrices of the individual intensity values. They also can be shown to diagonalize the cross-intensity covariance matrix. The

resulting values along the diagonal, δ_1 and δ_2, are the correlation coefficients between the maximum and minimum directional variances between the corresponding intensity values $I_1(\mathbf{p})$ and $I_2(\mathbf{p})$.

$$\mathbf{G}_1\Sigma_{I_1,I_2}\mathbf{G}_2^T = \mathbf{D} \qquad \mathbf{D} = \begin{pmatrix} \delta_1 & 0 \\ 0 & \delta_2 \end{pmatrix} \tag{5-33}$$

Taken together, equations (5-31), (5-32), and (5-33) specify the canonical correlations δ_1 and δ_2 and linear transformations $\mathbf{G}_1$ and $\mathbf{G}_2$ that incorporate the local image geometry. $\mathbf{G}_1$ and $\mathbf{G}_2$ can be chosen so that the δ_1 and δ_2 are both non-negative. Aside from the possibility of multiple roots (i.e., $\delta_1 = \delta_2$), δ_1 and δ_2 are unique. Under the ergodic assumption, $\mathbf{G}_1$ and $\mathbf{G}_2$ specify the directions for the multivalued covariance gauge.

5.8.2. Understanding Canonical Correlations of Multivalued Directional Statistics

Singular value decomposition of the directional covariance matrices of the individual intensity values determines the variance value λ_2 and the sampling direction $\mathbf{v}$ of the minimum influence of image geometry on local image statistics. In order to complete the picture of local probability distributions about a pixel in a 2-valued image, it is essential to understand how the two values covary in the minimum variance direction. The canonical correlation δ_2 provides that measure. Taken together, the minimum eigenvalues of the individual directional covariance matrices and δ_2 (the covariance between these two terms) can be used to describe a two-variable probability distribution.

δ_2 is easily rewritten as the correlation coefficient between I_1 sampled in the $\mathbf{v}_1$ direction and I_2 sampled in the $\mathbf{v}_2$ direction.

$$\delta_2 = \frac{\mu^{(2)}_{I_1,\mathbf{v}_1;I_2,\mathbf{v}_2}}{\sqrt{\lambda_{2,I_1}\,\lambda_{2,I_2}}} \tag{5-34}$$

In those locations where there is no directional bias, $\mathbf{K}_1$ and $\mathbf{K}_2$ become ambiguous, and isotropic sampling methods discussed in Chapter 4 are applicable. Near boundaries, the anisotropic spatial distribution of image intensities will align $\mathbf{K}_1$ and $\mathbf{K}_2$. If strong spatial anisotropy exists, $\mathbf{K}_1 \approx \mathbf{K}_2$, implying $\mathbf{u}_1 \approx \mathbf{u}_2$ and $\mathbf{v}_1 \approx \mathbf{v}_2$.

For 2-valued 2D images, the directional covariance matrix that reflects the minimum influence of image geometry on the local joint-intensity probability distribution is

$$\Lambda_2 = \begin{pmatrix} \lambda_{2,I_1} & \delta_2\sqrt{\lambda_{2,I_1}\,\lambda_{2,I_2}} \\ \delta_2\sqrt{\lambda_{2,I_1}\,\lambda_{2,I_2}} & \lambda_{2,I_2} \end{pmatrix} \tag{5-35}$$

This statistic, like the eigenvalue λ_2 in the context of scalar-valued images, can be used to normalize multivalued measurements. These normalized measurements are invariant with respect to linear functions of intensity applied to the separate intensity values that comprise the image. This normalization provides a common metric for

comparing the multiple parameters within pixels, enabling comparisons among incommensurable values.

5.9. Covariance between Image Intensity and Space

As suggested in the introduction to this chapter, an interesting development arises when canonical correlation analysis is applied to the way in which image space and image intensities covary. If the intensity values at location **p** are considered to be random variables and space is considered to be a random variable, a canonical correlation analysis can be performed on the resulting space of variables to orthogonalize their relationship. That is, canonical correlation analysis will yield a linear transformation that diagonalizes the covariance matrix describing the space of intensities and their spatial elements. It will also show the correlation between space, geometry, and noise for a given spatial location.

5.9.1. Directional Analysis of 2D Scalar Images

This section applies canonical correlation analysis directly to scalar images of two dimensions. Consider a two dimensional image I(**p**) with one intensity value per pixel. That is, I(**p**) is defined over $\mathbb{R}^1$ where $\mathbf{p} = (p_x, p_y) \in \mathbb{R}^2$. Following the formula for canonical correlation analysis described above and assigning the expectation operation to be a weighted spatial average, a relationship can be established between isotropic variance and the multiscale image gradient.

For each location $\mathbf{p}_0$ and its surrounding Gaussian-weighted neighborhood, analyze the correlations between pixels, viewing each pixel as a sample of a multivalued random variable. Consider the pixel location **p** = (x, y) to be a property of a pixel in the manner of a random variable. Also let the intensity value I(**p**) be treated as a random variable attached to the pixel. Given the neighborhood locus about $\mathbf{p}_0$, denote the mean of I(**p**) relative to $\mathbf{p}_0$ to be $\mu_I(\mathbf{p}_0, \sigma)$, a Gaussian-windowed average (with aperture σ) of intensities centered at $\mathbf{p}_0$. Treating space as a random variable, denote the mean of local space to be the central point $\mathbf{p}_0$.

Let Σ be the spatially weighted joint covariance matrix of I(**p**) and **p** about pixel $\mathbf{p}_0$. That is,

$$\Sigma = \begin{pmatrix} \Sigma_{I,I} & \Sigma_{I,p} \\ \Sigma_{p,I} & \Sigma_{p,p} \end{pmatrix} \tag{5-36}$$

where $\Sigma_{I,I}$ is 1×1 and corresponds to the local isotropic variance function. That is,

$$\Sigma_{I,I} = \mu_I^{(2)}(\mathbf{p} \mid \sigma) = G(\sigma, \mathbf{p}) \otimes (I(\mathbf{p}))^2 - (G(\sigma, \mathbf{p}) \otimes I(\mathbf{p}))^2 \tag{5-37}$$

This function is described in greater detail in Chapter 4. $\Sigma_{\mathbf{p},\mathbf{p}}$ is 2×2 and corresponds to

$$
\begin{aligned}
\Sigma_{\mathbf{p},\mathbf{p}} &= \begin{pmatrix} \left\langle (x-p_x)^2 \right\rangle & \left\langle (x-p_x)(y-p_y) \right\rangle \\ \left\langle (x-p_x)(y-p_y) \right\rangle & \left\langle (y-p_y)^2 \right\rangle \end{pmatrix} \\
&= \begin{pmatrix} \frac{1}{\sigma^2 2\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\tau^2 e^{-\frac{1}{2}\frac{\tau^2+\nu^2}{\sigma^2}}\,d\nu\,d\tau & \frac{1}{\sigma^2 2\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\tau\nu e^{-\frac{1}{2}\frac{\tau^2+\nu^2}{\sigma^2}}\,d\nu\,d\tau \\ \frac{1}{\sigma^2 2\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\tau\nu e^{-\frac{1}{2}\frac{\tau^2+\nu^2}{\sigma^2}}\,d\nu\,d\tau & \frac{1}{\sigma^2 2\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\nu^2 e^{-\frac{1}{2}\frac{\tau^2+\nu^2}{\sigma^2}}\,d\nu\,d\tau \end{pmatrix} \\
&= \begin{pmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}
\end{aligned}
\tag{5-38}
$$

$\Sigma_{\mathbf{I},\mathbf{p}}$ is 1×2, and $\Sigma_{\mathbf{I},\mathbf{p}} = (\Sigma_{\mathbf{p},\mathbf{I}})^T$.

$$
\begin{aligned}
\Sigma_{\mathbf{p},\mathbf{I}} &= \begin{pmatrix} \left\langle (xG(\sigma,x,y)) \otimes \left(\tilde{I}(x,y) - \mu_{\tilde{I}}(\mathbf{p})\right) \right\rangle \\ \left\langle (yG(\sigma,x,y)) \otimes \left(\tilde{I}(x,y) - \mu_{\tilde{I}}(\mathbf{p})\right) \right\rangle \end{pmatrix} \\
&= \begin{pmatrix} \frac{1}{\sigma^2 2\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} -\tau e^{-\frac{1}{2}\frac{\tau^2+\nu^2}{\sigma^2}} \left(\tilde{I}(p_x-\tau, p_y-\nu) - \mu_{\tilde{I}}(p_x,p_y)\right) d\nu\,d\tau \\ \frac{1}{\sigma^2 2\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} -\nu e^{-\frac{1}{2}\frac{\tau^2+\nu^2}{\sigma^2}} \left(\tilde{I}(p_x-\tau, p_y-\nu) - \mu_{\tilde{I}}(p_x,p_y)\right) d\nu\,d\tau \end{pmatrix} \\
&= \begin{pmatrix} \sigma^2 \frac{1}{\sigma^2 2\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} -\frac{\tau}{\sigma^2} e^{-\frac{1}{2}\frac{\tau^2+\nu^2}{\sigma^2}} \left(\tilde{I}(p_x-\tau, p_y-\nu) - \mu_{\tilde{I}}(\mathbf{p})\right) d\nu\,d\tau \\ \sigma^2 \frac{1}{\sigma^2 2\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} -\frac{\nu}{\sigma^2} e^{-\frac{1}{2}\frac{\tau^2+\nu^2}{\sigma^2}} \left(\tilde{I}(p_x-\tau, p_y-\nu) - \mu_{\tilde{I}}(\mathbf{p})\right) d\nu\,d\tau \end{pmatrix} \\
&= \begin{pmatrix} \sigma^2 \left(\frac{\partial}{\partial x} G(\sigma,\mathbf{p}) \otimes \tilde{I}(\mathbf{p})\right) \\ \sigma^2 \left(\frac{\partial}{\partial y} G(\sigma,\mathbf{p}) \otimes \tilde{I}(\mathbf{p})\right) \end{pmatrix} = \sigma^2 \left(\nabla\left(\tilde{I}(\mathbf{p}) \mid \sigma\right)\right)^T
\end{aligned}
\tag{5-39}
$$

The next step is to find δ, the root of

$$
\mathrm{Det}\begin{pmatrix} -\delta\Sigma_{\mathbf{I},\mathbf{I}} & \Sigma_{\mathbf{I},\mathbf{p}} \\ \Sigma_{\mathbf{p},\mathbf{I}} & -\delta\Sigma_{\mathbf{p},\mathbf{p}} \end{pmatrix} = 0 \tag{5-40}
$$

δ is the correlation between $I(\mathbf{p})$ and $\mathbf{p}$ in the Gaussian neighborhood about pixel $\mathbf{p}_0$. Computing (5-40) generates

$$\begin{aligned}
&\mathrm{Det}\begin{pmatrix} -\delta\sigma^2 & 0 \\ 0 & -\delta\sigma^2 \end{pmatrix} \mathrm{Det}\left(-\delta\mu_{\tilde{I}}^{(2)}(\mathbf{p}\mid\sigma) - \sigma^2 \left(\nabla\left(\tilde{I}(\mathbf{p})\mid\sigma\right)\right)^T \begin{pmatrix} -\frac{1}{\delta\sigma^2} & 0 \\ 0 & -\frac{1}{\delta\sigma^2} \end{pmatrix} \sigma^2 \left(\nabla\left(\tilde{I}(\mathbf{p})\mid\sigma\right)\right)\right) \\
&= \delta^2\sigma^4\left(-\delta\mu_{\tilde{I}}^{(2)}(\mathbf{p}\mid\sigma) + \tfrac{\sigma^2}{\delta}\left(\left(\tfrac{\partial}{\partial x}G(\sigma,\mathbf{p})\right)\otimes\tilde{I}(\mathbf{p})\right)^2 + \left(\left(\tfrac{\partial}{\partial y}G(\sigma,\mathbf{p})\right)\otimes\tilde{I}(\mathbf{p})\right)^2 \right) \\
&= \delta\sigma^6\left(\left(\left(\nabla\left(\tilde{I}(\mathbf{p})\mid\sigma\right)\right)\right)^2 \right) - \delta^3\sigma^4\mu_{\tilde{I}}^{(2)}(\mathbf{p}\mid\sigma)
\end{aligned} \tag{5-41}$$

Combining equations (5-40) and (5-41) yields

$$\begin{aligned}
\delta\sigma^6\left(\left(\nabla\left(\tilde{I}(\mathbf{p})|\sigma\right)\right)\right)^2 - \delta^3\sigma^4\mu_{\tilde{I}}^{(2)}(\mathbf{p}|\sigma) &= 0 \\
\delta\sigma^6\left(\left(\nabla\left(\tilde{I}(\mathbf{p})|\sigma\right)\right)\right)^2 &= \delta^3\sigma^4\mu_{\tilde{I}}^{(2)}(\mathbf{p}|\sigma) \\
\sigma^2\left(\left(\nabla\left(\tilde{I}(\mathbf{p})|\sigma\right)\right)\right)^2 &= \delta^2\mu_{\tilde{I}}^{(2)}(\mathbf{p}|\sigma) \\
\frac{\sigma^2\left(\left(\nabla\left(\tilde{I}(\mathbf{p})|\sigma\right)\right)\right)^2}{\mu_{\tilde{I}}^{(2)}(\mathbf{p}|\sigma)} &= \delta^2 \\
\frac{\sigma\left(\left(\nabla\left(\tilde{I}(\mathbf{p})|\sigma\right)\right)\right)}{\sqrt{\mu_{\tilde{I}}^{(2)}(\mathbf{p}|\sigma)}} &= \delta
\end{aligned} \tag{5-42}$$

Let G and **H** be invertible 1×1 and 2×2 matrices respectively, such that

$$G\Sigma_{I,I}G^T = \mathbf{I} \qquad \mathbf{H}\Sigma_{\mathbf{p},\mathbf{p}}\mathbf{H}^T = \mathbf{I} \qquad G\Sigma_{I,\mathbf{p}}\mathbf{H}^T = (D\ 0) \qquad D = \delta \tag{5-43}$$

In the current example, it is easily shown that

$$G = \frac{1}{\sqrt{\mu_{\tilde{I}}^{(2)}(\mathbf{p}\mid\sigma)}} \qquad \mathbf{H} = \begin{pmatrix} \frac{\cos\theta}{\sigma} & \frac{\sin\theta}{\sigma} \\ -\frac{\sin\theta}{\sigma} & \frac{\cos\theta}{\sigma} \end{pmatrix}$$

and
$$\frac{1}{\sqrt{\mu_{\tilde{I}}^{(2)}(\mathbf{p}\mid\sigma)}}\sigma^2\left(\nabla\left(\tilde{I}(\mathbf{p})\mid\sigma\right)\right)\begin{pmatrix} \frac{\cos\theta}{\sigma} & -\frac{\sin\theta}{\sigma} \\ \frac{\sin\theta}{\sigma} & \frac{\cos\theta}{\sigma} \end{pmatrix} = (\delta \quad 0) \tag{5-44}$$

where θ is the angle between the x axis and the gradient direction.

5.9.2. Canonical Correlation Analysis versus Differential Operators

This analysis shows the relationship between the scale space gradient direction, the scale space gradient magnitude, and the isotropic multiscale central moment. Specifically, for every position $\mathbf{p} = (p_x, p_y)$ in an image, there exist *canonical variables* throughout the neighborhood about **p**. That is, given image I(**p**), for any image position (x, y), there is a set of corresponding set of canonical variables A and **B** such that

$$A = \frac{I(x,y) - \mu_{\tilde{I}}(\mathbf{p} \mid \sigma)}{\sqrt{\mu^{(2)}_{\tilde{I}}(\mathbf{p} \mid \sigma)}} \qquad \mathbf{B} = \begin{pmatrix} \frac{\cos\theta}{\sigma} & \frac{\sin\theta}{\sigma} \\ -\frac{\sin\theta}{\sigma} & \frac{\cos\theta}{\sigma} \end{pmatrix} \begin{pmatrix} x - p_x \\ y - p_y \end{pmatrix} \tag{5-45}$$

Thus, any measurement made at position **p** uses a normalized multilocal coordinate frame, or normalized gauge coordinate system. It is also evident that these measurements are made relative to the mean and variance intensity of the local Gaussian neighborhood. The correlation between any A and **B** value is given by δ derived in equation (5-13).

These observations mirror similar methods to normalize and linearize scale space measurements. In particular, the use of gauge coordinates based on the image gradient has been proposed by Koenderink and refined for image analysis by ter Haar Romeny [ter Haar Romeny 1991ab]. Scale space normalization has been explored by many researchers including Eberly [Eberly 1994] and to a lesser extent Florack and ter Haar Romeny [Florack 1993][ter Haar Romeny 1993]. That these assertions can be derived through statistical methods lends credence to the relationship between the roots of this work in differential geometry and multiscale statistics.

These observations are included here for their relevance in relating the works of Chapter 3 and Chapter 4 in the setting of local directional image statistics. Further exploration into these relationships are necessary, especially in the realm of multivalued images. However, this dissertation is directed toward local isotropic and directional central moments of intensity, and a study of the canonical correlations among multivalued intensities and space is beyond the scope of this research.

5.10. Summary

This chapter has explored the derivation of directional multiscale image statistics. The derived statistical operators have been developed to comparable power as those based on differential geometry, with the added capability of being able to analyze multivalued images.

Joint second order central moments have been produced, and their properties discussed. A covariance matrix of these central moments was constructed and statistical tools adapted to analyze its properties. These statistics demonstrate invariance under

(1) Rotation

(2) Translation

(3) Scale

(4) Linear Functions of Intensity

Some example images were presented, demonstrating the application of these principles in a discrete image based setting. These multiscale directional central moments were generalized to images of multiple values and an analysis proposed using canonical correlation analysis. Finally, some observations on applying canonical correlation analysis directly to scalar images have been presented, indicating directions for future work.

5.A. Appendix: Singular Value Decomposition of a 2x2 Symmetric Matrix

These are some fundamentals of linear algebra used in this chapter. The analyses shown in sections 5.4. presume familiarity with singular value decomposition (or SVD), the diagonalization of symmetric matrices and the related extraction of their eigenvectors and eigenvalues. Eigenvalues of covariance matrices are required in earlier presentations.

Given a 2×2 symmetric matrix A of rank 2,

1) $A = A^T$
2) A is diagonalizable
3) The eigenvalues of A are real and they exist
4) The corresponding eigenvectors are (or can be chosen to be) orthogonal
5) The characteristic equation of A is

$$A = \begin{bmatrix} u & w \\ w & v \end{bmatrix} \quad \Rightarrow \quad \begin{array}{c} \lambda^2 - \text{Trace}(A)\lambda + \text{Deter min ant}(A) = 0 \\ \Leftrightarrow \lambda^2 - (u+v)\lambda + (uv - w^2) = 0 \end{array} \tag{5A-1}$$

6) Solving for the eigenvalues using the Quadratic formula

$$ax^2 + bx + c = 0 \quad \Rightarrow \quad x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \tag{5A-2}$$

$$\lambda_1 = \frac{(u+v) + \sqrt{(u+v)^2 - 4(uv + w^2)}}{2} \;,\; \lambda_2 = \frac{(u+v) - \sqrt{(u+v)^2 - 4(uv + w^2)}}{2} \tag{5A-3}$$

7) The eigenvectors are the solutions to the following equation:

$$\begin{array}{l} [\lambda_1 I - A]\begin{bmatrix} x_1 \\ y_1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ [\lambda_2 I - A]\begin{bmatrix} x_2 \\ y_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \end{array} \quad \Rightarrow \quad \begin{array}{cc} x_1(\lambda_1 - u) - y_1 w = 0 & y_1(\lambda_1 - v) - x_1 w = 0 \\ x_2(\lambda_2 - u) - y_2 w = 0 & y_2(\lambda_2 - v) - x_2 w = 0 \end{array} \tag{5A-4}$$

Solving the above equations yields the following values for the orthonormal basis for the corresponding eigenspace

$$\begin{array}{ll} x_1 = \dfrac{(\lambda_1 - v)}{\sqrt{w^2 + (\lambda_1 - v)^2}} & y_1 = \dfrac{w}{\sqrt{w^2 + (\lambda_1 - v)^2}} \\ x_2 = \dfrac{(\lambda_2 - v)}{\sqrt{w^2 + (\lambda_2 - v)^2}} & y_2 = \dfrac{w}{\sqrt{w^2 + (\lambda_2 - v)^2}} \end{array} \tag{5A-5}$$